

APPLICATION FOR UNITED STATES LETTERS PATENT

For

**METHOD AND SYSTEM TO RECOVER RESOURCES IN THE EVENT OF DATA  
BURST LOSS WITHIN WDM-BASED OPTICAL-SWITCHED NETWORKS**

Inventor(s):

**Shlomo Ovadia  
Christian Maciocco**

Prepared by:

**BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP  
12400 Wilshire Boulevard  
Los Angeles, CA 90025-1026  
(206) 292-8600**

Attorney's Docket No.: 42P17371

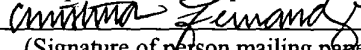
"Express Mail" mailing label number: EV320119210US

Date of Deposit: September 23, 2003

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Commissioner for Patents,  
Washington, D. C. 20231

CHRISTINA FERNANDEZ

(Typed or printed name of person mailing paper or fee)



(Signature of person mailing paper or fee)

September 23, 2003

(Date signed)

**METHOD AND SYSTEM TO RECOVER RESOURCES IN THE EVENT OF DATA  
BURST LOSS WITHIN WDM-BASED OPTICAL-SWITCHED NETWORKS**

**CROSS REFERENCE TO RELATED APPLICATIONS**

[0001] The present application is related to U.S. Patent Application No. 10/126,091, filed April 17, 2002; U.S. Patent Application No. 10/183,111, filed June 25, 2002; U. S. Patent Application No. 10/328,571, filed December 24, 2002; U.S. Patent Application No. 10/377,312 filed February 28, 2003; U.S. Patent Application No. 10/377,580 filed February 28, 2003; U.S. Patent Application No. 10/417,823 filed April 16, 2003; U.S. Patent Application No. 10/417,487 filed April 17, 2003; U.S. Patent Application No. (Attorney Docket No. 42P16183) filed May 19, 2003, U.S. Patent Application No. (Attorney Docket No. 42P16552) filed June 18, 2003, U.S. Patent Application No. (Attorney Docket No. 42P16847) filed June 14, 2003, and U.S. Patent Application No. (Attorney Docket No. 42P17373) filed August 6, 2003.

**FIELD OF THE INVENTION**

[0002] An embodiment of the present invention relates to optical networks in general; and, more specifically, to techniques for recovering resources in response to detection of resource unavailability within optical-switched networks.

**BACKGROUND INFORMATION**

[0003] Transmission bandwidth demands in telecommunication networks (*e.g.*, the Internet) appear to be ever increasing and solutions are being sought to support this bandwidth demand. One solution to this problem is to use fiber-optic networks, where wavelength-division-

multiplexing (WDM) technology is used to support the ever-growing demand in optical networks for higher data rates.

**[0004]** Conventional optical switched networks typically use wavelength routing techniques, which require that optical-electrical-optical (O-E-O) conversion of optical signals be done at the optical switching node. O-E-O conversion at each switching node in the optical network is not only very slow operation (typically about ten milliseconds), but it is very costly, power-consuming operation that potentially creates a traffic bottleneck for the optical switched network. In addition, the current optical switch technologies cannot efficiently support “bursty” traffic that is often experienced in packet communication applications (*e.g.*, the Internet).

**[0005]** A large enterprise data network can be implemented using many sub-networks. For example, a large enterprise network to support data traffic can be segmented into a large number of relatively small access networks, which are coupled to a number of local-area networks (LANs). The enterprise network is also coupled to metropolitan area networks (Optical MANs), which are in turn coupled to a large “backbone” wide area network (WAN). The optical MANs and WANs typically require a higher bandwidth than LANs in order to provide an adequate level of service demanded by their high-end users. However, as LAN speeds/bandwidth increase with improved technology, there is a need for increasing MAN/WAN speeds/bandwidth.

**[0006]** Recently, optical burst switching (OBS) scheme has emerged as a promising solution to support high-speed bursty data traffic over WDM optical networks. The OBS scheme offers a practical opportunity between the current optical circuit-switching and the emerging all optical packet switching technologies. It has been shown that under certain conditions, the OBS scheme achieves high-bandwidth utilization and class-of-service (CoS) by elimination of electronic

bottlenecks as a result of the O-E-O conversion occurring at switching nodes, and by using one-way end-to-end bandwidth reservation scheme with variable time slot duration provisioning scheduled by the ingress nodes. Optical switching fabrics are attractive because they offer at least one or more orders of magnitude lower power consumption with a smaller form factor than comparable O-E-O switches. However, most of the recently published work on OBS networks focuses on the next-generation backbone data networks (*i.e.* Internet-wide network) using high capacity (*i.e.*, 1 Tb/s) WDM switch fabrics with large number of input/output ports (*i.e.*, 256x256), optical channels (*i.e.*, 40 wavelengths), and requiring extensive buffering. Thus, these WDM switches tend to be complex, bulky, and very expensive to manufacture. In contrast, there is a growing demand to support a wide variety of bandwidth-demanding applications such as storage area networks (SANs) and multimedia multicast at a low cost for both LAN/WAN networks.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0007] Non-limiting and non-exhaustive embodiments of the present invention are described with reference to the following figures, wherein like reference numerals refer to like parts throughout the various views unless otherwise specified.

[0008] Figure 1 is a simplified block diagram illustrating a photonic burst-switched (PBS) network with variable time slot provisioning, according to one embodiment of the present invention.

[0009] Figure 2 is a simplified flow diagram illustrating the operation of a photonic burst-switched (PBS) network, according to one embodiment of the present invention.

[0010] Figure 3 is a block diagram illustrating a switching node module for use in a photonic burst-switched (PBS) network, according to one embodiment of the present invention.

[0011] Figure 4a is a diagram illustrating the format of an optical data burst for use in a photonic burst-switched network, according to one embodiment of the present invention.

[0012] Figure 4b is a diagram illustrating the format of an optical control burst for use in a photonic burst-switched network, according to one embodiment of the present invention.

[0013] Figure 5 is a flow diagram illustrating the operation of a switching node module, according to one embodiment of the present invention.

[0014] Figure 6 is a diagram illustrating a generalized multi-protocol label switching (GMPLS)-based architecture for a PBS network, according to one embodiment of the present invention.

[0015] Figure 7 is a diagram illustrating PBS optical burst flow between nodes in a PBS network, according to one embodiment of the present invention.

[0016] Figure 8 is a diagram illustrating generic PBS framing format for PBS optical bursts, according to one embodiment of the present invention.

[0017] Figure 9 is a diagram illustrating further details of the PBS framing format of Figure 8, according to one embodiment of the present invention.

[0018] Figure 10 is a schematic diagram illustrating an exemplary PBS network used to illustrate a resource recovery process, according to one embodiment of the present invention.

[0019] Figure 11 is a flowchart illustration operations performed in connection with a resource reservation and resource recovery process, according to one embodiment of the present invention.

[0020] Figure 12 is a diagram of an exemplary resource reservation table, according to one embodiment of the present invention.

[0021] Figure 13a is a diagram illustrating an extended PBS burst header that may be used in a resource cancellation PBS control burst, according to one embodiment of the present invention.

[0022] Figure 13b shows details of exemplary extended header data that may be stored in the PBS burst header of Figure 13a.

[0023] Figure 13c shows exemplary commands and corresponding command codes that may be stored in the command field of the extended PBS burst header of Figure 13a.

[0024] Figure 14 is a block diagram illustrating a GMPLS-based PBS label format, according to one embodiment of the present invention.

[0025] Figure 15 is a flowchart illustration the various operations performed in connection with the transmission and processing of control bursts, according to one embodiment of the present invention.

[0026] Figure 16 is a flowchart illustrating operations and logic performed during generation and processing of a resource cancellation message using a GMPLS-based PBS label, according to one embodiment of the present invention.

**[0027]** Figure 17 is a flowchart illustrating operations and logic performed during generation and processing of a resource cancellation message using a lightpath reservation identifier, according to one embodiment of the present invention.

**[0028]** Figure 18 is a schematic diagram of a PBS switching node architecture, according to one embodiment of the invention.

## DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0029] In the following detailed descriptions, embodiments of the invention are disclosed with reference to their use in a photonic burst-switched (PBS) network. A PBS network is a type of optical-switched network, typically comprising a high-speed hop and span-constrained network, such as an enterprise network. The term "photonic burst" is used herein to refer to statistically-multiplexed packets (*e.g.*, Internet protocol (IP) packets, Ethernet frames, Fibre Channel frames) having similar routing requirements. Although conceptually similar to backbone-based OBS networks, the design, operation, and performance requirements of these high-speed hop and span-constrained networks may be different. However, it will be understood that the teaching and principles disclosed herein may be applicable to other types of optical switched networks as well.

[0030] Figure 1 illustrates an exemplary photonic burst-switched (PBS) network 10 in which embodiments of the invention described herein may be implemented. A PBS network is a type of optical switched network. This embodiment of PBS network 10 includes local area networks (LANs) 13<sub>1</sub>-13<sub>N</sub> and a backbone optical WAN (not shown). In addition, this embodiment of PBS network 10 includes ingress nodes 15<sub>1</sub>-15<sub>M</sub>, switching nodes 17<sub>1</sub>-17<sub>L</sub>, and egress nodes 18<sub>1</sub>-18<sub>K</sub>. PBS network 10 can include other ingress, egress and switching nodes (not shown) that are interconnected with the switching nodes shown in Figure 1. The ingress and egress nodes are also referred to herein as edge nodes in that they logically reside at the edge of the PBS network. The edge nodes, in effect, provide an interface between the aforementioned "external" networks (*i.e.*, external to the PBS network) and the switching nodes of the PBS network. In this embodiment, the ingress, egress and switching nodes are implemented with



intelligent modules. This embodiment can be used, for example, as a metropolitan area network connecting a large number of LANs within the metropolitan area to a large optical backbone network.

**[0031]** In some embodiments, the ingress nodes perform optical-electrical (O-E) conversion of received optical signals, and include electronic memory to buffer the received signals until they are sent to the appropriate LAN/WAN. In addition, in some embodiments, the ingress nodes also perform electrical-optical (E-O) conversion of the received electrical signals before they are transmitted to switching nodes  $17_1$ - $17_M$  of PBS network 10.

**[0032]** Egress nodes are implemented with optical switching units or modules that are configured to receive optical signals from other nodes of PBS network 10 and route them to the optical WAN or other external networks. Egress nodes can also receive optical signals from the optical WAN or other external network and send them to the appropriate node of PBS network 10. In one embodiment, egress node  $18_1$  performs O-E-O conversion of received optical signals, and includes electronic memory to buffer received signals until they are sent to the appropriate node of PBS network 10 (or to the optical WAN).

**[0033]** Switching nodes  $17_1$ - $17_L$  are implemented with optical switching units or modules that are each configured to receive optical signals from other switching nodes and appropriately route the received optical signals to other switching nodes of PBS network 10. As is described below, the switching nodes perform O-E-O conversion of optical control bursts and network management control burst signals. In some embodiments, these optical control bursts and network management control bursts are propagated only on preselected wavelengths. The preselected wavelengths do not propagate optical “data” bursts (as opposed to control bursts and

network management control bursts) signals in such embodiments, even though the control bursts and network management control bursts may include necessary information for a particular group of optical data burst signals. The control and data information is transmitted on separate wavelengths in some embodiments (also referred to herein as out-of-band (OOB) signaling). In other embodiments, control and data information may be sent on the same wavelengths (also referred to herein as in-band (IB) signaling). In another embodiment, optical control bursts, network management control bursts, and optical data burst signals may be propagated on the same wavelength(s) using different encoding schemes such as different modulation formats, *etc.* In either approach, the optical control bursts and network management control bursts are sent asynchronously relative to its corresponding optical data burst signals. In still another embodiment, the optical control bursts and other control signals are propagated at different transmission rates as the optical data signals.

[0034] Although switching nodes  $17_1$ - $17_L$  may perform O-E-O conversion of the optical control signals, in this embodiment, the switching nodes do not perform O-E-O conversion of the optical data burst signals. Rather, switching nodes  $17_1$ - $17_L$  perform purely optical switching of the optical data burst signals. Thus, the switching nodes can include electronic circuitry to store and process the incoming optical control bursts and network management control bursts that were converted to an electronic form and use this information to configure photonic burst switch settings, and to properly route the optical data burst signals corresponding to the optical control bursts. The new control bursts, which replace the previous control bursts based on the new routing information, are converted to an optical control signal, and it is transmitted to the next switching or egress nodes. Embodiments of the switching nodes are described further below.

[0035] Elements of exemplary PBS network 10 are interconnected as follows. LANs  $13_1$ - $13_N$  are connected to corresponding ones of ingress nodes  $15_1$ - $15_M$ . Within PBS network 10, ingress nodes  $15_1$ - $15_M$  and egress nodes  $18_1$ - $18_K$  are connected to some of switching nodes  $17_1$ - $17_L$  via optical fibers. Switching nodes  $17_1$ - $17_L$  are also interconnected to each other via optical fibers in mesh architecture to form a relatively large number of lightpaths or optical links between the ingress nodes, and between ingress nodes  $15_1$ - $15_L$  and egress nodes  $18_1$ - $18_K$ . Ideally, there are more than one lightpath to connect the switching nodes  $17_1$ - $17_L$  to each of the endpoints of PBS network 10 (*i.e.*, the ingress nodes and egress nodes are endpoints within PBS network 10). Multiple lightpaths between switching nodes, ingress nodes, and egress nodes enable protection switching when one or more node fails, or can enable features such as primary and secondary route to destination.

[0036] As described below in conjunction with Figure 2, the ingress, egress and switching nodes of PBS network 10 are configured to send and/or receive optical control bursts, optical data burst, and other control signals that are wavelength multiplexed so as to propagate the optical control bursts and control labels on pre-selected wavelength(s) and optical data burst or payloads on different preselected wavelength(s). Still further, the edge nodes of PBS network 10 can send optical control burst signals while sending data out of PBS network 10 (either optical or electrical).

[0037] Figure 2 illustrates the operational flow of PBS network 10, according to one embodiment of the present invention. Referring to Figures 1 and 2, photonic burst switching network 10 operates as follows.

[0038] The process begins in a block 20, wherein PBS network 10 receives packets from LANs 13<sub>1</sub>-13<sub>N</sub>. In one embodiment, PBS network 10 receives IP packets at ingress nodes 15<sub>1</sub>-15<sub>M</sub>. The received packets can be in electronic form rather than in optical form, or received in optical form and then converted to electronic form. In this embodiment, the ingress nodes store the received packets electronically.

[0039] For clarity, the rest of the description of the operational flow of PBS network 10 focuses on the transport of information from ingress node 15<sub>1</sub> to egress node 18<sub>1</sub>. The transport of information from ingress nodes 15<sub>2</sub>-15<sub>M</sub> to egress node 18<sub>1</sub> (or other egress nodes) is substantially similar.

[0040] An optical burst label (*i.e.*, an optical control burst) and optical payload (*i.e.*, an optical data burst) is formed from the received packets, as depicted by a block 21. In one embodiment, ingress node 15<sub>1</sub> uses statistical multiplexing techniques to form the optical data burst from the received IP (Internet Protocol) packets stored in ingress node 15<sub>1</sub>. For example, packets received by ingress node 15<sub>1</sub> and having to pass through egress node 18<sub>1</sub> on their paths to a destination can be assembled into an optical data burst payload.

[0041] Next, in a block 22, Bandwidth on a specific optical channel and/or fiber is reserved to transport the optical data burst through PBS network 10. In one embodiment, ingress node 15<sub>1</sub> reserves a time slot (*i.e.*, a time slot of a TDM system) in an optical data signal path through PBS network 10. This time slot maybe fixed-time duration and/or variable-time duration with either uniform or non-uniform timing gaps between adjacent time slots. Further, in one embodiment, the bandwidth is reserved for a time period sufficient to transport the optical burst from the ingress node to the egress node. For example, in some embodiments, the ingress, egress, and

switching nodes maintain an updated list of all used and available time slots. The time slots can be allocated and distributed over multiple wavelengths and optical fibers. Thus, a reserved time slot (also referred to herein as a TDM channel), which in different embodiments may be of fixed-duration or variable-duration, may be in one wavelength of one fiber, and/or can be spread across multiple wavelengths and multiple optical fibers.

**[0042]** When an ingress and/or egress node reserves bandwidth or when bandwidth is released after an optical data burst is transported, a network controller (not shown) updates the list. In one embodiment, the network controller and the ingress or egress nodes perform this updating process using various burst or packet scheduling algorithms based on the available network resources and traffic patterns. The available variable-duration TDM channels, which are periodically broadcasted to all the ingress, switching, and egress nodes, are transmitted on the same wavelength as the optical control bursts or on a different common preselected wavelength throughout the optical network. The network controller function can reside in one of the ingress or egress nodes, or can be distributed across two or more ingress and/or egress nodes.

**[0043]** The optical control bursts, network management control labels, and optical data bursts are then transported through photonic burst switching network 10 in the reserved time slot or TDM channel, as depicted by a block 23. In one embodiment, ingress node 15<sub>1</sub> transmits the control burst to the next node along the optical label-switched path (OLSP) determined by the network controller. In this embodiment, the network controller uses a constraint-based routing protocol [e.g., multi-protocol label switching (MPLS)] over one or more wavelengths to determine the best available OLSP to the egress node.

**[0044]** In one embodiment, the control label (also referred to herein as a control burst) is transmitted asynchronously ahead of the photonic data burst and on a different wavelength and/or different fiber. The time offset between the control burst and the data burst allows each of the switching nodes to process the label and configure the photonic burst switches to appropriately switch before the arrival of the corresponding data burst. The term photonic burst switch is used herein to refer to fast optical switches that do not use O-E-O conversion.

**[0045]** In one embodiment, ingress node 15<sub>1</sub> then asynchronously transmits the optical data bursts to the switching nodes where the optical data bursts experience little or no time delay and no O-E-O conversion within each of the switching nodes. The optical control burst is always sent before the corresponding optical data burst is transmitted.

**[0046]** In some embodiments, the switching node may perform O-E-O conversion of the control bursts so that the node can extract and process the routing information contained in the label. Further, in some embodiments, the TDM channel is propagated in the same wavelengths that are used for propagating labels. Alternatively, the labels and payloads can be modulated on the same wavelength in the same optical fiber using different modulation formats. For example, optical labels can be transmitted using non-return-to-zero (NRZ) modulation format, while optical payloads are transmitted using return-to-zero (RZ) modulation format on the same wavelength. The optical burst is transmitted from one switching node to another switching node in a similar manner until the optical control and data bursts are terminated at egress node 18<sub>1</sub>.

**[0047]** The remaining set of operations pertains to egress node operations. Upon receiving the data burst, the egress node disassembles it to extract the IP packets or Ethernet frames in a block 24. In one embodiment, egress node 18<sub>1</sub> converts the optical data burst to electronic

signals that egress node 18<sub>1</sub> can process to recover the data segment of each of the packets. The operational flow at this point depends on whether the target network is an optical WAN or a LAN, as depicted by a decision block 25.

[0048] If the target network is an optical WAN, new optical label and payload signals are formed in a block 26. In this embodiment, egress node 18<sub>1</sub> prepares the new optical label and payload signals. The new optical label and payload are then transmitted to the target network (*i.e.*, WAN in this case) in a block 27. In this embodiment, egress node 18<sub>1</sub> includes an optical interface to transmit the optical label and payload to the optical WAN.

[0049] However, if in block 25 the target network is determined to be a LAN, the logic proceeds to a block 28. Accordingly, the extracted IP data packets or Ethernet frames are processed, combined with the corresponding IP labels, and then routed to the target network (*i.e.*, LAN in this case). In this embodiment, egress node 18<sub>1</sub> forms these new IP packets. The new IP packets are then transmitted to the target network (*i.e.*, LAN) as shown in block 29.

[0050] PBS network 10 can achieve increased bandwidth efficiency through the additional flexibility afforded by the TDM channels. Although this exemplary embodiment described above includes an optical MAN having ingress, switching and egress nodes to couple multiple LANs to an optical WAN backbone, in other embodiments the networks do not have to be LANs, optical MANs or WAN backbones. That is, PBS network 10 may include a number of relatively small networks that are coupled to a relatively larger network that in turn is coupled to a backbone network.

[0051] Figure 3 illustrates a module 17 for use as a switching node in photonic burst switching network 10 (Figure 1), according to one embodiment of the present invention. In this

embodiment, module 17 includes a set of optical wavelength division demultiplexers  $30_1-30_A$ , where  $A$  represents the number of input optical fibers used for propagating payloads, labels, and other network resources to the module. For example, in this embodiment, each input fiber could carry a set of  $C$  wavelengths (*i.e.*, WDM wavelengths), although in other embodiments the input optical fibers may carry differing numbers of wavelengths. Module 17 would also include a set of  $N \times N$  photonic burst switches  $32_1-32_B$ , where  $N$  is the number of input/output ports of each photonic burst switch. Thus, in this embodiment, the maximum number of wavelengths at each photonic burst switch is  $A \cdot C$ , where  $N \geq A \cdot C + 1$ . For embodiments in which  $N$  is greater than  $A \cdot C$ , the extra input/output ports can be used to loop back an optical signal for buffering.

**[0052]** Further, although photonic burst switches  $32_1-32_B$  are shown as separate units, they can be implemented as  $N \times N$  photonic burst switches using any suitable switch architecture. Module 17 also includes a set of optical wavelength division multiplexers  $34_1-34_A$ , a set of optical-to-electrical signal converters 36 (*e.g.*, photo-detectors), a control unit 37, and a set of electrical-to-optical signal converters 38 (*e.g.*, lasers). Control unit 37 may have one or more processors to execute software or firmware programs. Further details of control unit 37 are described below.

**[0053]** The elements of this embodiment of module 17 are interconnected as follows. Optical demultiplexers  $30_1-30_A$  are connected to a set of  $A$  input optical fibers that propagate input optical signals from other switching nodes of photonic burst switching network 10 (Figure 10). The output leads of the optical demultiplexers are connected to the set of  $B$  core optical switches  $32_1-32_B$  and to optical signal converter 36. For example, optical demultiplexer  $30_1$  has  $B$  output leads connected to input leads of the photonic burst



switches 32<sub>1</sub>-32<sub>B</sub> (*i.e.*, one output lead of optical demultiplexer 30<sub>1</sub> to one input lead of each photonic burst switch) and at least one output lead connected to optical signal converter 36.

[0054] The output leads of photonic burst switches 32<sub>1</sub>-32<sub>B</sub> are connected to optical multiplexers 34<sub>1</sub>-34<sub>A</sub>. For example, photonic burst switch 32<sub>1</sub> has *A* output leads connected to input leads of optical multiplexers 34<sub>1</sub>-34<sub>A</sub> (*i.e.*, one output lead of photonic burst switch 32<sub>1</sub> to one input lead of each optical multiplexer). Each optical multiplexer also an input lead connected to an output lead of electrical-to-optical signal converter 38. Control unit 37 has an input lead or port connected to the output lead or port of optical-to-electrical signal converter 36. The output leads of control unit 37 are connected to the control leads of photonic burst switches 32<sub>1</sub>-32<sub>B</sub> and electrical-to-optical signal converter 38. As described below in conjunction with the flow diagram of Figure 5, module 17 is used to receive and transmit optical control bursts, optical data bursts, and network management control bursts. In one embodiment, the optical data bursts and optical control bursts have transmission formats as shown in Figures 4A and 4B.

[0055] Figure 4A illustrates the format of an optical data burst for use in PBS network 10 (Figure 1), according to one embodiment of the present invention. In this embodiment, each optical data burst has a start guard band 40, an IP payload data segment 41, an IP header segment 42, a payload sync segment 43 (typically a small number of bits), and an end guard band 44 as shown in Figure 4A. In some embodiments, IP payload data segment 41 includes the statistically-multiplexed IP data packets or Ethernet frames used to form the burst. Although Figure 4A shows the payload as contiguous, module 17 transmits payloads in a TDM format. Further, in some embodiments the data burst can be segmented over multiple TDM channels. It

should be pointed out that in this embodiment the optical data bursts and optical control bursts have local significance only in PBS network 10, and may lose their significance at the optical WAN.

[0056] Figure 4B illustrates the format of an optical control burst for use in photonic burst switching network 10 (Figure 1), according to one embodiment of the present invention. In this embodiment, each optical control burst has a start guard band 46, an IP label data segment 47, a label sync segment 48 (typically a small number of bits), and an end guard band 49 as shown in Figure 4B. In this embodiment, label data segment 45 contains all the necessary routing and timing information of the IP packets to form the optical burst. Although Figure 4B shows the payload as contiguous, in this embodiment module 17 transmits labels in a TDM format.

[0057] In some embodiments, an optical network management control label (not shown) is also used in PBS network 10 (Figure 1). In such embodiments, each optical network management control burst includes: a start guard band similar to start guard band 46; a network management data segment similar to data segment 47; a network management sync segment (typically a small number of bits) similar to label sync segment 48; and an end guard band similar to end guard band 44. In this embodiment, network management data segment contains network management information needed to coordinate transmissions over the network. In some embodiments, the optical network management control burst is transmitted in a TDM format.

[0058] Figure 5 illustrates the operational flow of module 17 (Figure 3), according to one embodiment of the present invention. Referring to Figures 3 and 5, module 17 operates as follows.

[0059] Module 17 receives an optical signal with TDM label and data signals. In this embodiment, module 17 receives an optical control signal (*e.g.*, an optical control burst) and an optical data signal (*i.e.*, an optical data burst in this embodiment) at one or two of the optical demultiplexers. For example, the optical control signal may be modulated on a first wavelength of an optical signal received by optical demultiplexer 30<sub>A</sub>, while the optical data signal is modulated on a second wavelength of the optical signal received by optical demultiplexer 30<sub>A</sub>. In some embodiments, the optical control signal may be received by a first optical demultiplexer while the optical data signal is received by a second optical demultiplexer. Further, in some cases, only an optical control signal (*e.g.*, a network management control burst) is received. A block 51 represents this operation.

[0060] Module 17 converts the optical control signal into an electrical signal. In this embodiment, the optical control signal is the optical control burst signal, which is separated from the received optical data signal by the optical demultiplexer and sent to optical-to-electrical signal converter 36. In other embodiments, the optical control signal can be a network management control burst (previously described in conjunction with Figure 4B). Optical-to-electrical signal converter 36 converts the optical control signal into an electrical signal. For example, in one embodiment each portion of the TDM control signal is converted to an electrical signal. The electrical control signals received by control unit 37 are processed to form a new control signal. In this embodiment, control unit 37 stores and processes the information contained in the control signals. A block 53 represents this operation.

[0061] Module 17 then routes the optical data signals (*i.e.*, optical data burst in this embodiment) to one of optical multiplexers 34<sub>1</sub>-34<sub>A</sub>, based on routing information contained in

the control signal. In this embodiment, control unit 37 processes the control burst to extract the routing and timing information and sends appropriate PBS configuration signals to the set of  $B$  photonic burst switches  $32_1$ - $32_B$  to re-configure each of the photonic burst switches to switch the corresponding optical data bursts. A block 55 represents this operation.

[0062] Module 17 then converts the processed electrical control signal to a new optical control burst. In this embodiment, control unit 37 provides TDM channel alignment so that reconverted or new optical control bursts are generated in the desired wavelength and TDM time slot pattern. The new control burst may be modulated on a wavelength and/or time slot different from the wavelength and/or time slot of the control burst received in block 51. A block 57 represents this operation.

[0063] Module 17 then sends the optical control burst to the next switching node in the route. In this embodiment, electrical-to-optical signal generator 38 sends the new optical control burst to appropriate optical multiplexer of optical multiplexers  $34_1$ - $34_A$  to achieve the route. A block 59 represents this operation.

[0064] Figure 6 illustrates a GMPLS-based architecture for a PBS network, according to one embodiment of the present invention. Starting with the GMPLS suite of protocols, each of the GMPLS protocols can be modified or extended to support PBS operations and optical interfaces while still incorporating the GMPLS protocols' various traffic-engineering tasks. The integrated PBS layer architecture include PBS data services layer 60 on top of a PBS MAC layer 61, which is on top of a PBS photonics layer 62. It is well known that the GMPLS suite (indicated by a block 63 in Figure 6) includes a provisioning component 64, a signaling component 65, a routing component 66, a label management component 67, a link management component 68, and a

protection and restoration component 69. In some embodiments, these components are modified or have added extensions that support the PBS layers 60-62. Further, in this embodiment, GMPLS suite 63 is also extended to include an operation, administration, management and provisioning (OAM&P) component 70.

[0065] For example, signaling component 65 can include extensions specific to PBS networks such as, for example, burst start time, burst type, burst length, and burst priority, *etc.* Link management component 68 can be implemented based on the well known link management protocol (LMP) (that currently supports only SONET/SDH networks), with extensions added to support PBS networks. Protection and restoration component 69 can, for example, be modified to cover PBS networks.

[0066] Further, for example, label management component 67 can be modified to support a PBS control channel label space. In one embodiment, the label operations are performed after control channel signals are O-E converted. The ingress nodes of the PBS network act as label edge routers (LERs) while the switching nodes act as label switch routers (LSRs). An egress node acts as an egress LER substantially continuously providing all of the labels of the PBS network. This component can advantageously increase the speed of control channel context retrieval (by performing a pre-established label look-up instead of having to recover a full context).

[0067] Figure 7 illustrates PBS optical burst flow between nodes in an exemplary PBS network 700, according to one embodiment of the present invention. System 700 includes ingress node 710, a switching node 712, an egress node 714 and other nodes (egress, switching, and ingress that are not shown to avoid obscuring the description of the optical burst flow). In

this embodiment, the illustrated components of ingress, switching and egress nodes 710, 712 and 714 are implemented using machine-readable instructions that cause a machine (*e.g.*, a processor) to perform operations that allow the nodes to transfer information to and from other nodes in the PBS network. In this example, the lightpath for the optical burst flow is from ingress node 710, to switching node 712 and then to egress node 714.

[0068] Ingress node 710 includes an ingress PBS MAC layer component 720 having a data burst assembler 721, a data burst scheduler 722, an offset time manager 724, a control burst builder 726 and a burst framer 728. In one embodiment, data burst assembler 721 assembles the data bursts to be optically transmitted over PBS network 10 (Figure 1). In one embodiment, the size of the data burst is determined based on many different network parameters such as quality-of-service (QoS), number of available optical channels, the size of electronic buffering at the ingress nodes, the specific burst assembly algorithm, *etc.*

[0069] Data burst scheduler 722, in this embodiment, schedules the data burst transmission over PBS network 10 (Figure 1). In this embodiment, ingress PBS MAC layer component 710 generates a bandwidth request for insertion into the control burst associated with the data burst being formed. In one embodiment, data burst scheduler 722 also generates the schedule to include an offset time (from offset manager 724 described below) to allow for the various nodes in PBS network 10 to process the control burst before the associated data burst arrives.

[0070] In one embodiment, offset time manager 724 determines the offset time based on various network parameters such as, for example, the number of hops along the selected lightpath, the processing delay at each switching node, traffic loads for specific lightpaths, and class of service requirements.

[0071] Then control burst builder 726, in this embodiment, builds the control burst using information such as the required bandwidth, burst scheduling time, in-band or out-of-band signaling, burst destination address, data burst length, data burst channel wavelength, offset time, priorities, and the like.

[0072] Burst framer 728 frames the control and data bursts (using the framing format described below in conjunction with Figures 7-10 in some embodiments). Burst framer 728 then transmits the control burst over PBS network 10 via a physical optical interface (not shown), as indicated by an arrow 750. In this embodiment, the control burst is transmitted out of band (OOB) to switching node 712, as indicated by an optical control burst 756 and PBS TDM channel 757 in Figure 7. Burst framer 728 then transmits the data burst according to the schedule generated by burst scheduler 722 to switching node 712 over the PBS network via the physical optical interface, as indicated by an optical burst 758 and PBS TDM channel 759 in Figure 7. The time delay between optical bursts 756 (control burst) and 758 (data burst) is indicated as an  $OFFSET_1$  in Figure 7.

[0073] Switching node 712 includes a PBS switch controller 730 that has a control burst processing component 732, a burst framer/de-framer 734 and a hardware PBS switch (not shown).

[0074] In this example, optical control burst 756 is received via a physical optical interface (not shown) and optical switch (not shown) and converted to electrical signals (*i.e.*, O-E conversion). Control burst framer/de-framer 734 de-frames the control burst information and provides the control information to control burst processing component 732. Control burst

processing component 732 processes the information, determining the corresponding data burst's destination, bandwidth reservation, next control hop, control label swapping *etc.*

[0075] PBS switch controller component 730 uses some of this information to control and configure the optical switch (not shown) to switch the optical data burst at the appropriate time duration to the next node (*i.e.*, egress node 714 in this example) at the proper channel. In some embodiments, if the reserved bandwidth is not available, PBS switch controller component 730 can take appropriate action. For example, in one embodiment PBS switch controller 730 can: (a) determine a different lightpath to avoid the unavailable optical channel (*e.g.*, deflection routing); (b) delay the data bursts using integrated buffering elements within the PBS switch fabric such as fiber delay lines; (c) use a different optical channel (*e.g.* by using tunable wavelength converters); and/or (d) drop only the coetaneous data bursts. Some embodiments of PBS switch controller component 730 may also send a negative acknowledgment message back to ingress node 710 to re-transmit the dropped burst.

[0076] However, if the bandwidth can be found and reserved for the data burst, PBS switch controller component 730 provides appropriate control of the hardware PBS switch (not shown). In addition, PBS switch controller component 730 generates a new control burst based on the updated reserved bandwidth from control burst processing component 732 and the available PBS network resources. Control burst framer/de-framer 734 then frames the re-built control burst, which is then optically transmitted to egress node 714 via the physical optical interface (not shown) and the optical switch (not shown), as indicated by PBS TDM channel 764 and an optical control burst 766 in Figure 7.



[0077] Subsequently, when the optical data burst corresponding to the received/processed control burst is received by switching node 712, the hardware PBS switch is already configured to switch the optical data burst to egress node 714. In other situations, switching node 712 can switch the optical data burst to a different node (*e.g.*, another switching node not shown in Figure 7). The optical data burst from ingress node 710 is then switched to egress node 714, as indicated by PBS TDM channel 767 and an optical data burst 758A. In this embodiment, optical data burst 758A is simply optical data burst 758 re-routed by the hardware PBS switch (not shown), but possibly transmitted in a different TDM channel. The time delay between optical control burst 766 and optical data burst 758A is indicated by an  $OFFSET_2$  in Figure 7, which is smaller than  $OFFSET_1$  due, for example, to processing delay and other timing errors in switching node 712.

[0078] Egress node 714 includes a PBS MAC component 740 that has a data demultiplexer 742, a data burst re-assembler 744, a control burst processing component 746, and a data burst de-framer 748.

[0079] Egress node 714 receives the optical control burst as indicated by an arrow 770 in Figure 7. Burst de-framer 748 receives and de-frames the control burst via a physical O-E interface (not shown). In this embodiment, control burst processing component 746 processes the de-framed control burst to extract the pertinent control/address information.

[0080] After the control burst is received, egress node 714 receives the data burst(s) corresponding to the received control burst, as indicated by an arrow 772 in Figure 7. In this example, egress node 714 receives the optical data burst after a delay of  $OFFSET_2$ , relative to the end of the control burst. In a manner similar to that described above for received control bursts,

burst de-framer 748 receives and de-frames the data burst. Data burst re-assembler 744 then processes the de-framed data burst to extract the data (and to re-assemble the data if the data burst was a fragmented data burst). Data de-multiplexer 742 then appropriately de-multiplexes the extracted data for transmission to the appropriate destination (which can be a network other than the PBS network).

[0081] Figure 8 illustrates a generic PBS framing format 800 for PBS optical bursts, according to one embodiment of the present invention. Generic PBS frame 800 includes a PBS generic burst header 802 and a PBS burst payload 804 (which can be either a control burst or a data burst). Figure 8 also includes an expanded view of PBS generic burst header 802 and PBS burst payload 804.

[0082] PBS generic burst header 802 is common for all types of PBS bursts and includes a version number (VN) field 810, a payload type (PT) field 812, a control priority (CP) field 814, an in-band signaling (IB) field 816, a label present (LP) field 818, a header error correction (HEC) present (HP) field 819, a burst length field 822, and a burst ID field 824. In some embodiments, PBS generic burst header also includes a reserved field 820 and a HEC field 826. Specific field sizes and definitions are described below for framing format having 32-bit words; however, in other embodiments, the sizes, order and definitions can be different.

[0083] In this embodiment, PBS generic burst header 802 is a 4-word header. The first header word includes VN field 810, PT field 812, CP field 814, IB field 816 and LP field 818. VN field 810 in this exemplary embodiment is a 4-bit field (*e.g.*, bits 0-3) defining the version number of the PBS Framing format being used to frame the PBS burst. In this embodiment, VN

field 810 is defined as the first 4-bits of the first word, but in other embodiments, it need not be the first 4-bits, in the first word, or limited to 4-bits.

**[0084]** PT field 812 is a 4-bit field (bits 4-7) that defines the payload type. For example, binary "0000" may indicate that the PBS burst is a data burst, while binary "0001" indicates that the PBS burst is a control burst, and binary "0010" indicates that the PBS burst is a management burst. In this embodiment, PT field 812 is defined as the second 4-bits of the first word, but in other embodiments, it need not be the second 4-bits, in the first word, or limited to 4-bits.

**[0085]** CP field 814 is a 2-bit field (bits 8-9) that defines the burst's priority. For example, binary "00" may indicate a normal priority while binary "01" indicates a high priority. In this embodiment, PT field 812 is defined bits 8 and 9 of the first word, but in other embodiments, it need not be bits 8 and 9, in the first word, or limited to 2-bits.

**[0086]** IB field 816 is a one-bit field (bit 10) that indicates whether the PBS control burst is being signaled in-band or OOB. For example, binary "0" may indicate OOB signaling while binary "1" indicates in-band signaling. In this embodiment, IB field 816 is defined as bit 10 of the first word, but in other embodiments, it need not be bit 10, in the first word, or limited to one-bit.

**[0087]** LP field 818 is a one-bit field (bit 11) used to indicate whether a label has been established for the lightpath carrying this header. In this embodiment, LP field 818 is defined as bit 11 of the first word, but in other embodiments, it need not be bit 11, in the first word, or limited to one-bit.

[0088] HP field 819 is a one-bit (bit 12) used to indicate whether header error correction is being used in this control burst. In this embodiment, HP field 819 is defined as bit 12 of the first word, but in other embodiments, it need not be bit 12, in the first word, or limited to one-bit. The unused bits (bits 13-31) form field(s) 820 that are currently unused and reserved for future use.

[0089] The second word in PBS generic burst header 802, in this embodiment, contains PBS burst length field 822, which is used to store a binary value equal to the length the number of bytes in PBS burst payload 804. In this embodiment, the PBS burst length field is 32-bits. In other embodiments, PBS burst length field 822 need not be in the second word and is not limited to 32-bits.

[0090] In this embodiment, the third word in PBS generic burst header 802 contains PBS burst ID field 824, which is used to store an identification number for this burst. In this embodiment, PBS burst ID field 824 is 32-bits generated by the ingress node (*e.g.*, ingress node 710 in Figure 7). In other embodiments, PBS burst ID field 824 need not be in the third word and is not limited to 32-bits.

[0091] The fourth word in PBS generic burst header 802, in this embodiment, contains generic burst header HEC field 826, which is used to store an error correction word. In this embodiment, generic burst header HEC field 826 is 32-bits generated using any suitable known error correction technique. In other embodiments, generic burst header HEC field 826 need not be in the fourth word and is not limited to 32-bits. As indicated in Figure 8, generic burst header HEC field 826 is optional in that if error correction is not used, the field may be filled

with all zeros. In other embodiments, generic burst header HEC field 826 is not included in PBS generic burst header 802.

[0092] PBS burst payload 804 is common for all types of PBS bursts and includes a PBS specific payload header field 832, a payload field 834, and a payload frame check sequence (FCS) field 836.

[0093] In this exemplary embodiment, PBS specific payload header 832 is the first part (*i.e.*, one or more words) of PBS burst payload 804. Specific payload header field 832 for a control burst is described below in more detail in conjunction with Figure 9. Similarly, specific payload field 832 for a data burst is described below in conjunction with Figure 9. Typically, specific payload header field 832 includes one or more fields for information related to a data burst, which can be either this burst itself or contained in another burst associated with this burst (*i.e.*, when this burst is a control burst).

[0094] Payload data field 834, in this embodiment, is the next portion of PBS burst payload 804. In some embodiments, control bursts have no payload data, so this field may be omitted or contain all zeros. For data bursts, payload data field 834 may be relatively large (*e.g.*, containing multiple IP packets or Ethernet frames).

[0095] Payload FCS field 836, in this embodiment, is the next portion of PBS burst payload. In this embodiment, payload FCS field 836 is a one-word field (*i.e.*, 32-bits) used in error detection and/or correction. As in indicated in Figure 8, payload FCS field 836 is optional in that if error detection/correction is not used, the field may be filled with all zeros. In other embodiments, payload FCS field 836 is not included in PBS burst payload 804.

[0096] Figure 9 illustrates a PBS optical control burst framing format 900, according to one embodiment of the present invention. To help improve clarity, Figure 9 includes the expanded views of PBS generic burst header 802 and PBS burst payload 804 (previously described in conjunction with Figure 8), with a further expansion of PBS payload header field 832 (described below) when part of a control burst. In this example, the PT field is set to "01" to indicate that the burst is a control burst. The CP field is set to "0" to indicate that the burst has normal priority. The IB field is set to "0" to indicate that the burst is using OOB signaling. The LP field is set to "0" to indicate that there is no label for this control burst.

[0097] In this exemplary embodiment of a PBS control burst, PBS payload header field 832 includes: a PBS control length field 902; an extended header (EH) field 906; an address type (AT) field 908; a payload FCS present (PH) field 910; a control channel wavelength field 920; a data channel wavelength field 922; a PBS label field 924; a PBS data burst length field 926; a PBS data burst start time field 930; a PBS data burst time-to-live (TTL) field 932; a data burst priority field 934; a PBS data burst destination address field 938; and an optional extended header field 940.

[0098] In this embodiment, the first word of PBS payload header 832 includes PBS control length field 902, which is used for storing the length of the control header in bytes. In this embodiment, PBS control length field 902 is a 16-bit field (bits 0-15) calculated by control burst builder 726 (Figure 7) or control burst processor 732 (Figure 7). In other embodiments, PBS control length field 902 need not be the first 16-bits, in the first word, or limited to 16-bits. A reserved field 904 (bits 16-27) is included in PBS payload header 832 in this embodiment. In other embodiments, these bits may be used for other field(s).

[0099] The first word of PBS payload header 832 also includes EH field 906, which is used in this embodiment to indicate whether an extended header is present in the burst. In this embodiment, EH field 906 is a 1-bit field (bit 28). In other embodiments, EH field 906 need not be bit 28, or in the first word.

[00100] The first word of PBS payload header 832 also includes AT field 908, which is used in this embodiment to indicate the address type of the associated PBS data burst's destination. For example, the address type may be an IP address (*e.g.*, IPv4, IPv6), a network service access point (NSAP) address, an Ethernet address or other type of address. In this embodiment, AT field 908 is a 2-bit field (bits 29-30). In other embodiments, AT field 908 need not be bits 17-18, in the first word, or limited to 2-bits.

[00101] In this embodiment, the first word of PBS payload header 832 also includes PH field 910, which is used to indicate whether a payload FCS is present in the burst. In this embodiment, PH field 910 is a 1-bit field (bit 31). In other embodiments, EH field 906 need not be bit 16, or in the first word.

[00102] The second word of PBS payload header 832, in this embodiment, includes control channel wavelength field 920, which is used to indicate a WDM wavelength in which the control burst is supposed to be modulated. In this embodiment, control channel wavelength field 920 is a 16-bit field (bits 0-15). In other embodiments, control channel wavelength field 920 need not be bits 0-15, in the second word, or limited to 16-bits.

[00103] In this embodiment, the second word of PBS payload header 832 also includes data channel wavelength field 922, which is used to indicate a WDM wavelength in which the data burst is to be modulated. In this embodiment, data channel wavelength field 922 is a 16-bit field

(bits 16-31). In other embodiments, data channel wavelength field 922 need not be bits 16-31, in the second word, or limited to 16-bits.

**[00104]** A third word of PBS payload header 832 includes PBS label field 924, which is used in this embodiment to store the label (if any) for the lightpath being used by the burst. In this embodiment, the label is a 32-bit word generated by label management component 67 (Figure 6). In other embodiments, PBS label field 924 need not be the third word, or limited to 32-bits.

**[00105]** A fourth word of PBS payload header 832 includes PBS data burst length field 926. In this embodiment, the PBS data burst length is a 32-bit word. In other embodiments, PBS data burst length field 926 need not be the fourth word, or limited to 32-bits.

**[00106]** A fifth word of PBS payload header 832 includes PBS data burst start time field 930. In this embodiment, the PBS data burst start time is a 32-bit word, generated by burst scheduler 722 (Figure 7). In other embodiments, PBS data burst start time field 930 need not be the fifth word, or limited to 32-bits.

**[00107]** A sixth word of PBS payload header 832 includes PBS data TTL field 932. In this embodiment, PBS data TTL field 932 is a 16-bit (bits 0-15) field, generated by ingress PBS MAC component 720 (Figure 7). For example, in one embodiment, burst scheduler 722 (Figure 7) of ingress PBS MAC component 720 can generate the TTL value. In other embodiments, PBS data TTL field 932 need not be bits 0-15, in the sixth word, or limited to 16-bits.

**[00108]** The sixth word of PBS payload header 832 also includes data burst priority field 932. In this embodiment, data burst priority field 932 is an 8-bit field (bits 16-23), generated by



ingress PBS MAC component 720 (Figure 7). For example, in one embodiment, burst scheduler 722 (Figure 7) of ingress PBS MAC component 720 can generate the data burst priority value. In other embodiments, data burst priority field 932 need not be bits 16-23, in the sixth word, or limited to 8-bits. Further, in this embodiment, the sixth word of PBS payload header 832 includes a reserved field 936 (bits 24-31) which can be used in the future for other field(s).

**[00109]** A seventh word of PBS payload header 832 also includes PBS data burst destination address field 938. In this embodiment, PBS data burst destination address field 938 is variable length field, shown as a single 32-bit word for clarity. In other embodiments, PBS data burst destination address field 938 need not be limited to 32-bits. The actual length of the address may vary, depending on the address type as indicated in AT field 908.

**[00110]** An eight word of PBS payload header 832 can include extended header field 940. This header can be used to hold other header data that may be used in the future. When this header is used, EH field 906 is set to 1. In this embodiment, payload data field 834 and payload FCS field 836 have been described above.

**[00111]** In accordance with further aspects of embodiments of the invention, mechanisms are now disclosed for recovering node (switching or end node) resources in response to detection of a resource failure. For example, under PBS operations, a lightpath comprising a plurality of lightpath segments is reserved for a given variable-duration timeslot via corresponding control bursts. Each switching node along the route (as identified by an incoming lightpath segment received at that switching node) maintains a reservation table containing reservation data indicating how it is to switch incoming and outgoing data corresponding to currently-reserved

timeslots. If a switching node failures (*e.g.*, a fiber gets cut or disconnected or bandwidth is determined to be unavailable due to traffic constraints, etc), the lightpath cannot be completed for the current data burst. As a result, any network resources (*i.e.* both external to a node and resources provided internally by a node) reserved along the lightpath by a corresponding control burst no longer will be used for routing subsequently-sent data bursts. Under conventional approaches, the use of these resources for the reserved timeslots would simply be lost. However, under embodiments of the resource recovery mechanism, information is passed to appropriate switching nodes to inform those nodes that the resources will not be used, and thus are freed up to accept new reservations spanning from the initially-reserved timeslots.

**[00112]** An exemplary lightpath reservation and corresponding resource recovery is illustrated in Figure 10. Figure 10 shows an exemplary PBS network 1000, including PBS switching nodes 1, 2, 3, 4, 5, 6, and 7. The PBS switching nodes are linked via various fibers, including fiber links 1002, 1004, 1006, 1008, 1010, 1012, 1014, 1016, 1018, 1020, 1022, 1024, 1026, 1028, 1030, 1032, and 1033. PBS network 1000 further includes edge nodes A, B, C, D, E, and F, which are linked to corresponding switching nodes via fiber links 1034, 1036, 1038, 1040, 1042, and 1044. The edge nodes A, B, C, D, E, and F provide ingress and/or egress points to PBS network 1000, enabling external networks 1046, 1048, 1050, 1052, 1054, and 1056 to route data between each other via PBS network 1000, wherein respective links 1058, 1060, 1062, 1064, 1066, and 1068 are used to connect the external networks to a respective edge node. From the viewpoint of each external network, they only can "see" the edge node to which they are connected, and are aware of other external networks that may be reached via PBS network 1000. In effect, the rest of PBS network 1000 appears as a "black box" to the external networks, and, in

fact, the external networks do not need to be aware of the internal infrastructure of a PBS network.

[00113] Figure 11 shows a flowchart illustrating the operations that are performed during resource reservation and cancellation in response to detected resource unavailability in accordance with one embodiment. The process begins in a block 1300, in which an ingress node generates a control burst in response to a network access request, and the control burst is then routed between the ingress node and intermediate switching nodes until it reaches the egress node to which the destination network is coupled to set up resource reservations along the lightpath. For example, in the illustrated example of Figure 10, it is desired to send traffic comprising a data burst from external network 1040 (*i.e.*, the source) to external network 1056 (*i.e.*, the destination). Thus, the ingress node will be edge node A, while the egress node will be edge node F. Accordingly, a control burst is generated at edge node A having a format discussed above with reference to Figures 8 and 9, and sent out to reserve resources along a route build by concatenating multiple lightpath segments to form a lightpath between the ingress and egress edge nodes A and F. An exemplary route (lightpath) shown in Figure 10 is depicted using a dash-dot-dash line format, and includes lightpath segments 1034, 1004, 1016, 1028, 1032, and 1044, which are coupled between edge node A, switching node 1, switching node 3, switching node 5, switching node 6, switching node 7 and edge node F, respectively.

[00114] As the control burst is processed at each switching node, a reservation comprising a scheduled allocation of bandwidth for specified input and output lightpath segments at specified wavelengths for a specified timeslot is made, as depicted by a block 1102. In one embodiment, reservation data are stored in a reservation table 1200, as shown in Figure 12. Reservation

table 1200 includes a plurality of columns (fields) in which data are stored in rows, wherein data for a given row is called a "record." The columns include an optional Key column 1202, an Input Fiber Port column 1204, an Input Wavelength column 1206, an optional Input Lightpath Segment ID column 1208, an Output Fiber Port column 1210, and Output Wavelength column 1212, an optional Output Lightpath Segment ID column 1214, a Start Time column 1216, and End Time column 1218, and an optional Status column 1220

[00115] In general, Key column 1202 is used to store a unique identifier for each record, enabling quick retrieval of records and guaranteeing record uniqueness. In one embodiment, this unique identifier comprises the PBS burst ID stored in PBS burst field 824 of the control burst.

[00116] Data relating to incoming link parameters are stored in Input Fiber Port column 1204, Input Wavelength column 1206, and, optionally, Input Lightpath Segment ID column 1208, while data relating to outgoing link parameters are stored in Output Fiber Port column 1210, Output Wavelength column 1212, and, optionally, Output Lightpath Segment ID column 1214. Each switching node is coupled to two or more fiber links via respective I/O ports. For example, the exemplary data in reservation table 1200 corresponds to switching node 3, which includes six network input/output (I/O) ports, depicted as encircled numbers 1-6. The value in Input Fiber Port column 1204 identifies the I/O port at which data is received by the particular switching node maintaining the reservation table, while the value in Output Fiber Port column 1210 identifies the I/O port via which data is transmitted. As an option, input and output fiber routing data may be stored through reference to input and output lightpath segments, in lieu of or in addition to specifying I/O ports. Accordingly, in one embodiment, data identifying the input and

output lightpath segments is stored in Input Lightpath Segment ID column 1208 and Output Lightpath Segment ID column 1214, respectively.

[00117] As discussed above, a given lightpath segment may support concurrent data streams that are transmitted using different wavelengths. Accordingly, data in Input Wavelength column 1206 is used to identify the wavelength incoming data is to be transmitted at for a given reservation record, while data in Output Wavelength column 1212 is used to identify the wavelength outgoing data is to be transmitted at.

[00118] Routing paths through each switching node are reserved for a respective timeslot of variable duration based on appropriate information contained in the control burst. Typically, the time slot will be defined by a start time and an end time, with corresponding data stored in Start Time column 1216 and End time column 1218. In one embodiment, the start time comprises an offset from the time at which the control burst is processed by the switching node. Optionally, a start time may be specified by PBS data burst start time field 930.

[00119] The end time for a given reservation is stored in End Time 1218. The end time, in effect, will comprise an offset from the start time, wherein the offset is selected such that the full data burst may be successfully transmitted from source to destination without reserving any unnecessary extra time or bandwidth. Typically, a control burst will reserve a resource timeslot for an amount of time varying from microseconds to milliseconds, although longer timeslots may also be reserved. For simplicity, the time data shown Start Time column 1116 and End Time column 1118 only reflect the millisecond level. The length of the reservation request, as specified in PBS data burst length field 926, will be determined as a function of the data payload (*i.e.*, size of the payload in bytes) and the transmission bandwidth (*e.g.*, 1 gigabits/sec, 10

gigabits/sec, *etc.*) For example, a payload of 1 megabits would require 1 millisecond to be transferred over a 1 gigabit/s Ethernet (1 GbE) link.

[00120] Optional Status column 1120 is used for providing status information relating to the reservation. For example, a binary value may be used to indicate whether a reservation is valid or invalid (*i.e.*, cancelled). Optionally, a multi-bit value may be used to specifying one of a plurality of reservation statuses.

[00121] Continuing with the flowchart of Figure 11, the remaining operations concern detection of resource non-availability and corresponding release of resource reservations. First, in a block 1104, an unavailable switching node resource is detected. Non-availability of a resource will generally result from resource constraints due to traffic contention or a switching node or fiber link failure. For instance, in the illustrated example it is presumed that a traffic fault is detected that indicates lightpath segment 1016 is unavailable for transmitting data bursts corresponding to the resource reservation made in response to processing the control burst. When the data burst arrives at each switching node along the reserved lightpath, it may encounter traffic contention. In other words, two different incoming data bursts at two different ports at switching node have the same switching node output port destination (for switches that support multiple concurrent wavelength transmissions) at the same time. If, for example, the incoming data bursts are classified according to their priority, then the simplest way to resolve this contention is to drop the lower priority incoming data bursts while transmitting the higher priority data bursts. However, the control burst of the dropped data burst continues to reserve the necessary bandwidth with the PBS switch configurations on the subsequent upstream switching nodes along its lightpath until it is terminated at the destination egress node. Consequently, this

leads to a wasted reserved bandwidth, since the upstream reserved switch configurations for the dropped data burst cannot be used by other data bursts until the reserved bandwidth is released. Therefore, the overall network throughput is decreased with increased end-to-end latency.

[00122] Embodiments of the invention address this problem through a bandwidth (*i.e.*, resource) cancellation mechanism that is implemented via an extension to the PBS signaling protocol discussed above. In particular, the extended PBS signaling protocol has the ability to signal to switching nodes along the reserved lightpath (either upstream and/or downstream nodes) that a specific switching node has dropped a data burst due to resource constraints or switch/link failures, and that the corresponding resource reservations along the selected lightpath can now be released and made available to other data burst reservation requests. The mechanism is initiated in a block 1106, wherein a “Resource Cancellation Message” (RCM), which has a similar format to the control burst, is generated at the switching node at which the non-availability is detected.

[00123] In one embodiment, the mechanism employs a variant of the control burst format shown in Figures 8 and 9, wherein a resource cancellation control burst is propagated along the lightpath route in response to a resource unavailable or failure detection. The control interface unit within the switching node where the data burst was dropped (or otherwise detecting a switch or link failure) generates a RCM that is transmitted to appropriate nodes along the selected lightpath up to the destination egress node. Depending on the particular implementation and/or type of failure, the resource cancellation message may be sent to upstream (from the failure point forward to the destination egress node) switching nodes and/or downstream (from the failure point backward to the source ingress node) switching nodes.

[00124] In one embodiment, the Resource Cancellation Message is stored in the extended header field 826 of a control burst. For example, Figure 13a shows the format of an extended header field 826A that may contain data so as to function as a Resource Cancellation Message. The extended header includes a command field 1300, a Reserved (R) field 1302, a PAD field 1304, a Length field 1306, and Extended Header data 1308. The Command field 1300 comprises a 12-bit field that identifies the Command carried by the Extended Header, *e.g.*, a command indicating a "Bandwidth Cancellation" operation. The Reserved field 1302 is a 1-bit field containing a reserved bit. The PAD field 1304 comprises a 3-bit field that identifies the number of padding bytes that may be necessary to pad the last word of the extended header field to form a 32-bit word.

[00125] The length bit comprises a 16-bit field that contains the length, in words (*i.e.*, 32-bits), including the Command/Length word of the Extended Header. The minimum length will be "1", *i.e.*, a Command field only for commands that do not require any associated data. The Extended Header data field 1308 is a variable-length field that may contain various types of information. The field may employ up to 3 bytes of padding

[00126] The intermediate optical switching node where the resource contention occurs has all the necessary optical burst state information pertinent to the data burst that was just discarded or is being discarded due to the detected resource constraint or switch/link failure. This information is employed to build a PBS Control Burst Frame that is generated at this node that will be propagated along the same hop-by-hop lightpath used by the control frame that reserved the bandwidth initially. This time, however, the control unit of the switching node also fills up the



Extended Header field of the control frame. For example, an exemplary set of extended header data includes the following values:

Command: 0x001 (Bandwidth Cancellation)

PAD: 0 (for IPv4 address type, might have a value for other type of addressed based on the AT field of the header)

Length: 1+n (1 for header + n for extended data)

Extended Header Data:

- Address of node where the failure occurred
- Type of Failure (i.e., traffic contention, fiber link failure, switching node failure, etc.)
- Label stack: All the labels (a, b,..., etc.) used along the reserved lightpath

**[00127]** Figure 13b shows further details of exemplary data that may be contained in the extended header data field 1308. In addition to identifying the address of the node that has failed or is otherwise unavailable, the data may identify a type of failure such as traffic contention, fiber link failure, switching node failure, etc. As described below, the labels are used for routing resource reservations. In one embodiment, resources are released at the label level rather than the node level, unless a command value indicates that all resources are to be released.

**[00128]** In one embodiment, the command field 1300 contains a value or code that defines how resource cancellation is to be implemented. For example, exemplary actions and command codes are shown in Figure 13c. The simplest action is to cancel the reservation for the resource at the affected node. The most complex action is to cancel resource reservations for all nodes along the lightpath. Other action options include canceling resource reservations for upstream or downstream nodes, as described below.

[00129] Additional information for the resource cancellation control burst may be derived from the control burst that was previously sent to establish the resource reservation. This includes the PBS burst ID stored in PBS burst ID field 824, which may be used to uniquely identify the lightpath for which resource reservations are made.

[00130] Once the resource cancellation control burst is generated, it is routed upstream and/or downstream along the lightpath so that it is received and processed at corresponding switching and edge nodes, as indicated in a block 1108. In one embodiment, the mechanism for routing the resource cancellation control burst is similar to that employed for routing a "normal" control burst. In general, data is extracted at each switching node is used to determine the "next hop" in the lightpath chain. For example, in one embodiment, data from reservation table 1200 is extracted to determine the next hop. When the PBS burst ID is stored in Key column 1102, corresponding next-hop routing information for both upstream and downstream nodes can be easily extracted. First, the reservation record is retrieved based on the PBS burst ID value. Once retrieved, the next upstream hop corresponds to the switching or edge node connected to the fiber link coupled to the output fiber port specified by the value in Output Fiber Port column 1210 or identified by the value in Output Lightpath Segment ID column 1214. Similarly, the next downstream hop corresponds to the switching or edge node connected to the fiber link couple to the input fiber port (as specified by the value in Input Fiber Port column 1204) or identified by the value in Input Lightpath Segment ID column 1208.

[00131] In one embodiment, resource cancellation messages containing GMPLS-based labels are employed to route the resource cancellation message between nodes. For example, label management component 67 can be modified to support a PBS control channel message space. In

one embodiment, the label operations are performed after control channel signals are O-E converted. The ingress nodes of the PBS network act as label edge routers (LERs) while the switching nodes act as label switch routers (LSRs). An egress node acts as an egress LER substantially continuously providing all of the labels of the PBS network. An ingress node can propose a label to be used on the lightpath segment it is connected to, but the downstream node will be the deciding one in selecting the label value, potentially rejecting the proposed label and selecting its own label. A label list can also be proposed by a node to its downstream node. This component can advantageously increase the speed of control channel context retrieval (by performing a pre-established label look-up instead of having to recover a full context). Further details of the label usage and processing are described below in connection with Figure 16.

[00132] Returning to the flowchart of Figure 13, in a block 1310, processing of the resource cancellation control burst is performed, resulting in cancellation of the corresponding resource reservations. For example, a resource reservation may be cancelled by deleting (*i.e.*, removing) the record specified by the PBS burst ID, or marking the record as invalid via a change to the value in Status column 1220. As each switching node considers existing reservations when determining whether to accept a reservation request, canceling the resource reservation has the effect of releasing the resource for subsequent use during the reserved timeslot.

[00133] As discussed above, in one embodiment the resource reservation cancellation process is facilitated through use of a GMPLS-based label scheme. The signaling of PBS labels for lightpath set-up, tear down, and maintenance is done through an extension of IETF (internet engineering task force) resource reservation protocol-traffic engineering (RSVP-TE). More

information on GMPLS signaling with RSVP-TE extensions can be found at <http://www.ietf.org/rfc/rfc3473.txt>.

[00134] The PBS label, which identifies the data burst input fiber, wavelength, and lightpath segment, channel spacing, is used on the control path to enable one to make soft reservation request of the network resources (through corresponding RESV messages). If the request is fulfilled (through the PATH message), each switching node along the selected lightpath commits the requested resources, and the lightpath is established with the appropriate segment-to-segment labels. Each switching node is responsible for updating the initial PBS label through the signaling mechanism, indicating to the previous switching node the label for its lightpath segment. If the request cannot be fulfilled or an error occurred, a message describing the condition is sent back to the originator to take the appropriate action (i.e., select another lightpath characteristics). Thus, the implementation of the PBS label through signaling enables an MPLS type efficient lookup for the control burst processing. This processing improvement of the control burst at each switching node reduces the required offset time between the control and data bursts, resulting in an improved PBS network throughput and reduced end-to-end latency.

[00135] In one embodiment, the label signaling scheme reduces the PBS offset time by reducing the amount of time it takes to process a signaled lightpath. This is achieved by extending the GMPLS model to identify each lightpath segment within the PBS network using a unique label defined in a PBS label space. The use of a PBS label speeds up the PBS control burst processing by allowing the control interface unit within the PBS switching node, which processes the control burst, to lookup relevant physical routing information and other relevant processing state based on the label information used to perform a fast and efficient lookup.

Thus, each PBS switching node has access in one lookup operation to the following relevant information, among others: 1) the address of the next hop to send the control burst to; 2) information about the outgoing fiber and wavelength; 3) label to use on the next segment if working in a label-based mode; and 4) data needed to update the scheduling requirement for the specific input port and wavelength.

**[00136]** An exemplary GMPLS-based PBS label format 1400 is shown in Figure 14 with its corresponding fields. In the illustrated embodiment, PBS label 1400 comprises five fields, including an input fiber port field 1402, a input wavelength field 1404, a lightpath segment ID field 1406, a channel spacing ( $\Delta$ ) field 1408, and a reserved field 1410. The input fiber port field 1402 comprises an 8-bit field that specifies the input fiber port of the data channel identified by the label (which itself is carried on the control wavelength). The input wavelength field 1704 comprises a 32-bit field that describes the input data wavelength used on the input fiber port specified by input fiber port field 1402. In one embodiment, the input wavelength is represented using IEEE (Institute of Electrical and Electronic Engineers) standard 754 for single precision floating-point format. The 32-bit word is divided into a 1-bit sign indicator  $S$ , an 8-bit biased exponent  $e$ , and a 23-bit fraction. The lightpath segment ID field 1406 comprises a 16-bit field that describes the lightpath segment ID on a specific wavelength and a fiber cable. Lightpath segment ID's are predefined values that are determined based on the PBS network topology. The channel spacing field 1408 comprises a 4-bit field used for identifying the channel spacing (*i.e.*, separation between adjacent channels).

**[00137]** The transmitted PBS control bursts, which are processed electronically by the PBS Network processor (NP), undergo the following operations: With reference to the flowchart of

Figure 15, the process begins in a block 1500, wherein the control burst is de-framed, classified according to its priority, and the bandwidth reservation information is processed. If an optical flow has been signaled and established this flow label is used to lookup the relevant information. Next, in a block 1502, the PBS switch configuration settings for the reserved bandwidth on the selected wavelength at a specific time is either confirmed or denied. If confirmed, the process proceeds; if denied, a new reservation request process is initiated.

[00138] In a block 1504, PBS contention resolution is processed in case of PBS switch configuration conflict. One of the three possible contention resolution schemes, namely FDL-based buffering, tunable wavelength converters, and deflection routing can be selected. If none of these schemes are available, the incoming data bursts are dropped until the PBS switch becomes available and a negative acknowledgement message is sent to the ingress node to retransmit. A new control burst is generated in a block 1506, based on updated network resources retrieved from the resource manager, and scheduled for transmission. The new control burst is then framed and placed in the output queue for transmission to the next node in a block 1508.

[00139] With reference to the flowchart of Figure 16, further details of the operations of blocks 1106, 1108, and 1110 in accordance with one embodiment that employs the foregoing PBS labels and associated data are illustrated. The process begins in a block 1600 in which input labels corresponding to the unavailable resource are identified at the detecting node. For example, Columns 1204, 1206 and 1208 of resource reservation table 1200 contains data extracted from input labels during the resource reservation process. (It is noted that input wavelength column 1206 shows a numerical input wavelength value for illustrative purposes.

The input wavelength data contained in input wavelength field 1404 and channel spacing field 1408 may also be stored in separate columns.). In general, determination of unavailable resources will identify input lightpath segment and/or input fiber port. In some instances, the unavailable resource may pertain to a particular input wavelength for a given lightpath segment.

[00140] Once the labels are identified, corresponding resource reservation records are retrieved from resource reservation table 1200 in a block 1602. The resource reservation tables are then grouped based on the next hop(s) identified by the outgoing label data in a block 1604. For example, for upstream next hops, the next hop information may be identified by the output fiber port and/or the output lightpath segment ID data contained in the retrieved records, while for downstream next hops, the next hop information may be identified by the input fiber port and/or the input lightpath segment ID data.

[00141] Next, in a block 1606, an initial Resource Cancellation Message (RCM), identifying relevant labels (for a given group), are generated for each next hop. These messages are then sent to the next hops. Sending the data can be accomplished by broadcasting the message on the applicable output fiber port (such that it is received by the next hop node), or sending the data to the next-hop address, which can be retrieved based on local network topology information stored at the node. For example, the node may store information that correlates input and output fiber ports with corresponding address. Activities for the detecting node are completed in a block 1608 by canceling the resource reservations (records) containing the identified labels.

[00142] Subsequent processing operations performed at each next hop are shown in the lower portion of the flowchart delineated by start and end loop blocks 1610 and 1620. These operations are similar to those performed at the detecting node. First, in a block 1612, resource

reservation records are retrieved that include label data corresponding to the labels identified in the resource cancellation message. As before, the resource reservation records are then grouped by corresponding next hops in a block 1614. An updated resource cancellation message identifying the relevant input or output labels for each next hop are then generated and sent in a block 1616. The resource reservation records containing the label data are then cancelled in a block 1618. This process is repeated until the final nodes along the lightpaths (e.g., an ingress or egress node) are reached.

**[00143]** The flowchart of Figure 17 includes further details of the operations of blocks 1106, 1108, and 1110 in accordance with one embodiment that employs the foregoing label data in combination with the control burst ID data. In this embodiment, the operations of blocks 1600, 1602, and 1604 are performed in the same manner as discussed above; thus, at a block 1706 resource reservation records corresponding to the unavailable resource are retrieved and grouped by next hop. In block 1706, a single resource cancellation message identifying the lightpaths for which resources are to be released is generated and sent to each next hop, as applicable. In one embodiment, the lightpath is identified by the control burst ID value contained in Key column 1202. Since reservations for a given lightpath are made in response to the same control burst, the control burst ID for the control burst may be used to link the resource reservation records stored at the nodes along the lightpath together. The resource reservation records containing the identified lightpaths (e.g., lightpath ID's) are then cancelled at the detecting node, releasing the corresponding resources.

**[00144]** The operations performed at each next hop are shown in the lower portion of the flowchart delineated by start and end loop blocks 1710 and 1720. These operations are similar to



those performed at the detecting node. First, in a block 1712, resource reservation records are retrieved from the resource reservation table at the current node based on the lightpath ID's. The resource reservation records are then grouped by corresponding next hops in a block 1714. The resource cancellation message is then sent to the next hop(s), as applicable, in a block 1716. The resource reservation records containing the lightpath ID's are then cancelled in a block 1718. This process is repeated until the final nodes along the lightpaths (e.g., an ingress or egress node) are reached.

#### Switching Node Architecture

**[00145]** A simplified block diagram 1800 of a PBS switching node architecture in accordance with one embodiment is shown in Figure 18. The intelligent switching node architecture is logically divided into control plane components and data plane. The control plane includes a control unit 37 employing a network processor (NP) 1802, coupled to glue logic 1804 and a control processor (CPU) 1806 that runs software components to perform the GMPLS control operations 1808 disclosed herein. Network processor 1802 is also coupled to one or more banks of SDRAM (synchronous dynamic random access memory) memory 1810. The data plane architecture comprises a non-blocking optical switch fabric including a PBS 32 coupled optical multiplexers 1812, de-multiplexers 1814, and optical transceivers (as depicted by a receive (Rx) block 1816 and a transmit (Tx) block 1818).

**[00146]** The burst assembly and framing, burst scheduling and control, which are part of the PBS MAC layer and related tasks are performed by network processor 1802. Network processors are very powerful processors with flexible micro-architecture that are suitable to support wide-range of packet processing tasks, including classification, metering, policing,

congestion avoidance, and traffic scheduling. For example, the Intel® IXP2800 NP, which has 16 microengines, can support the execution of up to 1493 microengines instructions per packet at packet rate of 15 million packets per second for 10 GbE and a clock rate of 1.4 GHz.

**[00147]** In one embodiment, the optical switch fabric has strictly non-blocking space-division architecture with fast ( $< 100$  ns) switching times and with limited number of input/output ports (*e.g.*,  $\approx 8 \times 8$ ,  $12 \times 12$ ). Each of the incoming or outgoing fiber links typically carries only one data burst wavelength. The switch fabric, which has no or limited optical buffering fabric, performs statistical burst switching within a variable-duration time slot between the input and output ports. The PBS network can operate with a relatively small number of control wavelengths ( $\lambda'_0, \lambda_0$ ), since they can be shared among many data wavelengths. Furthermore, the PBS switch fabric can also operate with a single wavelength and multiple fiber; however, further details of this implementation are not disclosed herein.

**[00148]** The control bursts can be sent either in-band (IB) or out of band (OOB) on separate optical channels. For the OOB case, the optical data bursts are statistically switched at a given wavelength between the input and output ports within a variable time duration by the PBS fabric based on the reserved switch configuration as set dynamically by network processor 1802. NP 1802 is responsible to extract the routing information from the incoming control bursts, providing fix-duration reservation of the PBS switch resources for the requested data bursts, and forming the new outgoing control bursts for the next PBS switching node on the path to the egress node. In addition, the network processor provides overall PBS network management functionality based on then extended GMPLS framework discussed above. For the IB case, both the control and data bursts are transmitted to the PBS switch fabric and control interface unit.

However, NP 1802 ignores the incoming data bursts based on the burst payload header information. Similarly, the transmitted control bursts are ignored at the PBS fabric since the switch configuration has not been reserved for them. One advantage of this approach is that it is simpler and cost less to implement since it reduces the number of required wavelengths.

**[00149]** Embodiments of method and apparatus for implementing a photonic burst switching network are described herein. In the above description, numerous specific details are set forth to provide a thorough understanding of embodiments of the invention. One skilled in the relevant art will recognize, however, that embodiments of the invention can be practiced without one or more of the specific details, or with other methods, components, materials, *etc.* In other instances, well-known structures, materials, or operations are not shown or described in detail to avoid obscuring this description.

**[00150]** Reference throughout this specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, the appearances of the phrases “in one embodiment” or “in an embodiment” in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be combined in any suitable optical manner in one or more embodiments.

**[00151]** Thus, embodiments of this invention may be used as or to support software program executed upon some form of processing core (such as the CPU of a computer or a processor of a module) or otherwise implemented or realized upon or within a machine-readable medium. A machine-readable medium includes any mechanism for storing or transmitting information in a

form readable by a machine (*e.g.*, a computer). For example, a machine-readable medium can include such as a read only memory (ROM); a random access memory (RAM); a magnetic disk storage media; an optical storage media; and a flash memory device, *etc.* In addition, a machine-readable medium can include propagated signals such as electrical, optical, acoustical or other form of propagated signals (*e.g.*, carrier waves, infrared signals, digital signals, *etc.*).

**[00152]** In the foregoing specification, embodiments of the invention have been described. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.